

device determines that one or more weights of the height histogram is greater than a threshold and generates one or more plane hypotheses corresponding to the one or more respective height ranges. For example, with reference to FIG. 7, the device generates a first plane hypothesis corresponding to the first height bin 701 and a second plane hypothesis corresponding to the fourth height bin 704. The first plane hypothesis describes a first horizontal plane (e.g., normal to the gravity vector) at a height corresponding to the first height bin 701 (e.g., its center height) and the second plane hypothesis describes a second horizontal plane at a height corresponding to the fourth height bin 704.

[0071] In various implementations, the device detects one or more peaks in the height histogram and generates one or more horizontal plane hypotheses corresponding to the peaks. In various implementations, the device applies a non-maxima suppression algorithm to the height histogram in detecting the one or more peaks. For example, in various implementations, the device filters the height histogram.

[0072] Whereas FIG. 6 describes detection of horizontal planes, the method 600 can also be used to detect any of a set of planes with one degree of freedom by generating a histogram indicative of the number of points of the plurality of points associated with each of a plurality of locations along an axis. For example, the method 600 can be used to detect the presence of planes that are normal to the line-of-sight of the electronic device by generating a histogram indicative of the number of points of the plurality of points associated with each of a plurality of depths along an axis parallel to the line-of-sight.

[0073] FIG. 8 is a flowchart representation of a method 800 of generating a vertical plane hypothesis in accordance with some implementations. In various implementations, the method 800 is performed by a device with one or more processors, non-transitory memory, and a scene camera (e.g., the HMD 120 FIG. 3). In some implementations, the method 800 is performed by processing logic, including hardware, firmware, software, or a combination thereof. In some implementations, the method 800 is performed by a processor executing instructions (e.g., code) stored in a non-transitory computer-readable medium (e.g., a memory). Briefly, in some circumstances, the method 800 includes: obtaining a point cloud of a scene, detecting one or more lines describing at least a threshold number of points of the point cloud (as projected to a common height), and generating one or more vertical plane hypotheses based on the one or more lines.

[0074] The method 800 begins, in block 802, with the device obtaining a point cloud of a scene including a plurality of points in a gravity-aligned coordinate system. The point cloud may be obtained using a depth sensor, VIO, other computer vision techniques, or the like. In various implementations, each of the plurality of points is associated with three coordinates in the gravity-aligned coordinate system, one of the three coordinates (e.g., a “z-coordinate” or a “height coordinate”) corresponding to a height of the point. In various implementations, each of the plurality of points is further associated with an uncertainty, e.g., an uncertainty (or, similarly, a confidence) in the position of the point on the surface in the scene. In various implementations, a point in the point cloud is associated with metadata, such as a color, texture, reflectance, or transmittance of the point on the surface in the scene.

[0075] In various implementations, the plurality of points includes one or more unpaired points. In various implementations, the plurality of points includes one or more sets of paired points. Each of the set of paired points defines a line having a length.

[0076] The method 800 continues, at block 804, with the device generating, based on the plurality of points, a line-space image, each pixel of the line-space image corresponding to a two-parameter representation of a respective line at a common height in the gravity-aligned coordinate system and each pixel having a pixel value. In various implementations, the device projects the plurality of points to a common height in the gravity-aligned coordinate system by setting the height coordinate of each point to zero or simply ignoring the height coordinate of each point to obtain a two-dimensional (2D) point map. Thus, in various implementations, the device detects the one or more lines by processing the plurality of points of the point cloud without the one of the three coordinates corresponding to the height of the point.

[0077] FIG. 9A illustrates an example 2D point map 901 based on the point cloud represented in FIG. 5. The 2D point map 901 includes a first cluster of points 916 corresponding to points of the point cloud corresponding to the side wall 406. The 2D point map 901 includes a second cluster of points 917 corresponding to points of the point cloud corresponding to the back wall 407. The 2D point map 901 includes a third cluster of points 909 corresponding to points of the point cloud corresponding to the table 409. The 2D point map 901 includes other points 918, spread out, corresponding to points of the point cloud corresponding to the floor 408.

[0078] The 2D point map 901 includes unpaired points, such as a point 931 corresponding to unpaired point 431 of FIG. 5, and paired points, such as paired points 932A-932B corresponding to paired points 432A-432B of FIG. 5.

[0079] FIG. 9A illustrates a first line 910a, a second line 910b, and a third line 910c overlaid on the 2D point map 901. The first line 910a represents all of the points of the first cluster of points 916, one point from the second cluster of points 917, and one other point 918. The second line 910b represents a few of the other points 918 and a few points of the second cluster of points 917. The third line 910c represents all of the points of the second cluster of points 917.

[0080] In various implementations, a line is considered to represent a point in the 2D point map 901 when the point is less than a threshold distance from the line. In various implementations, the distance is an L1-distance or an L2-distance. In various implementations, when a point is associated with an uncertainty, a line is considered to represent the point when the likelihood that the point lies on the line is above a threshold.

[0081] As noted above, each pixel of the line-space image corresponding to a two-parameter representation of a respective line at the common height in the gravity-aligned coordinate system. For example, a pixel at location (u,v) corresponds to a line along a ground of the gravity-aligned coordinate system represented by the parameters u and v. In various implementations, the line can be represented in slope-intercept form (e.g., $y=mx+b$), where m and b are functions of u and v. In various implementations, the line can be represented in distance-angle or Hesse normal form (e.g., $r=x \cos(\theta)+y \sin(\theta)$), where r and θ are functions of u and v.